Super Resolution in MRI: How far can we go?

H Joan Coward[a] and Gareth J Barker[a*]

[a]Centre for Neuroimaging Sciences, Institute of Psychiatry, King's College London

Abstract. Super Resolution (SR) in MRI has been described as a technique that allows an increase of resolution without a loss in the Signal to Noise Ratio (SNR), essentially providing "free" resolution. This study explores this claim and investigates the limits to SR in the presence of noise. We show that the amplification of noise due to the SR algorithm can outweigh the signal benefit, leading to SNR reductions that could counter the motivation for using the SR approach. Thresholds for successful SR approaches are determined through simulation, which are lower than used previously in the literature. We also assess the initial resolving power of different slice profile shapes, without any further post-processing, and show that for rectangular slice profiles, it can be far higher than the slice width suggests. This finding has practical implications for structural MRI and could be used to provide "free" SNR.

1 Introduction

Super Resolution (SR) is a term used primarily for image processing methods that achieve an improvement in image resolution from lower resolution input images [1]. It is a topical area of research in many types of imaging, including satellite and surveillance, where the expense or practicality of increasing resolution by conventional means may not be feasible. SR has been applied to MRI with some controversy, most notably the first attempt, which applied it in the in-plane direction [2]. This approach was questioned because the MR signal is inherently bandlimited, hence cannot contain any information at frequencies higher than the limiting pixel size [3, 4]. However the application of SR in the through-plane direction offers greater potential for MRI applications, and has already been used for fMRI and cardiac MRI as a method enabling an increase in slice resolution without a loss in Signal to Noise Ratio (SNR) [5–7].

The application of SR in the through-plane direction would consist of the following acquisition and processing. Acquire data with a slice thickness of st mm (normally measured by the FWHM of the slice profile), at a slice separation of ss mm, $st > ss$. This effectively corresponds to overlapping slices, where the Slice Overlap Ratio, SOR = st/ss. Then apply a SR algorithm to recover the ideal resolution ss mm; this would consist of a deconvolution of the slice profile from the data. A perfect SR reconstruction would therefore result in higher resolution post-processed data compared to the input data, allowing an increase in resolution without loss of SNR associated with acquiring directly at SOR=1. However this scenario is in the case of noise-free images, for which the SR algorithm is able to deblur the data and restore the resolution to ss mm, with no noise amplification. Unfortunately all deconvolution methods are sensitive to the presence of noise, which causes non-uniqueness and instability in the solution [8].

We hypothesise that, in the presence of noise, there is a limiting SOR at which deconvolution can successfully recover the true high-resolution image, without amplifying the noise to such an extent that there is no further SNR benefit to the SR approach. The aim of this work is to investigate the limiting SOR levels under different noise conditions and slice profiles. This is done initially using simulations of a Resolution Test Phantom, and then illustrated using MRI data. The results are discussed in the context of SOR values used in previous implementations of SR in MRI.

2 Methods

All simulations are carried out using MATLAB (The MathWorks, Inc).

2.1 Resolution Test Phantom

Object and Image Simulation. The Resolution Test Phantom is composed of two objects, each of width ss, their centres separated by $2ss$. To simulate the imaging process in the MR scanner, the Test Phantom is convolved over a high-resolution grid, $0.01ss$, with the slice profile, and then discretely sampled at ss intervals to form the image. Images are simulated for either a Gaussian or rectangular slice profiles, and a range of SOR values (1, 1.1, ..., 1.8). Random Gaussian noise is then added, with the standard deviation, n_{pre}, scaled to the Noise/Signal (1/SNR) ratio; a range of noise values are investigated between 2% and 20%, chosen to simulate a range of applications from high SNR (e.g. structural) to low SNR (e.g. Diffusion Weighted Imaging, DWI).

*Correspondence to Joan.Coward@iop.kcl.ac.uk

Resolution Assessment. The phantom allows a quantitative assessment of resolution by calculating the signal modulation between the two objects. A modulation of 100% corresponds to a perfect reconstruction, and the objects are said to be resolved if the modulation is equal or greater than 50% [9]. For this study, the modulation is calculated as $(\min(x_1, x_2) - x_0)/\min(x_1, x_2)$, where x_1 and x_2 are the image object intensities, and x_0 is the image intensity between the objects. The initial modulation for the simulated images is calculated prior to deconvolution.

SR algorithm. A regularised filter with a Laplacian smoothing constraint is used for the deconvolution, the performance of which can be altered by varying its regularisation parameter, λ, the Lagrange Multiplier. The value of λ determines either the dominance of noise-propagation error (i.e. noisy solutions), or the dominance of approximation error, (i.e. blurry images) [8]. Since the imaging model begins as a continuous model, and is then discretised by assuming a zero-hold interpolation, the PSF used in the deconvolution is a discretely sampled convolution of the continuous slice profile with a rectangular function representing the zero-hold.

Our SR algorithm optimises the deconvolution by selecting the maximum $\lambda = \lambda_{opt}$, for which the deconvolved image can resolve the objects. This is done on an iterative basis, beginning with a high λ, and gradually reducing it until the objects can just be resolved. Once λ_{opt} has been found, the post SR algorithm noise in the deconvolved image is then calculated by estimating the standard deviation in a background region, n_{post}. This algorithm is then repeated 1000 times with different noise realisations, to build up a distribution of results.

SNR assessment. The SNR change, relative to SOR=1 and no SR algorithm, can then be calculated from $\Delta\text{SNR} = \Delta\text{signal}/\Delta\text{noise} = \text{SOR}/(n_{post}/n_{pre})$. An SNR change greater than 1 indicates that the signal change from the thicker slice is greater than the corresponding noise amplification from the SR algorithm (necessary to achieve the true resolution of ss mm). However if the SNR change is below 1, then the noise amplification exceeds the signal increase, implying that there is no net SNR benefit from using the SR algorithm over acquiring directly at the ss resolution (assuming this is technologically possible). If the SNR change is greater than 1, then this can be directly traded for isotropic resolution according to the following formula, $\Delta x_{iso} = 1 - 1/\sqrt[3]{\Delta\text{SNR}}$, which indicates the improvement in resolution for the same "SNR-quality" of data.

2.2 MR Test Data

A slice from a 1.1mm^3 SPGR data set is used as the input high-resolution image. The imaging process is simulated as described earlier for the Resolution Phantom, for $ss = 2$mm and the same range of SOR values. Deconvolution is then carried out using the regularisation parameter values, λ_{opt}, as determined previously.

3 Results and Discussion

3.1 Resolution Test Phantom

Initial Resolution Assessment. Figure 1 shows histograms of the number of successful resolutions prior to deconvolution. The histogram on the left shows the results for a Gaussian slice profile, indicating that with no noise present, the objects can be successfully resolved with an SOR=1.2, without the need for any SR algorithm. The corresponding SOR value for the rectangular slice profile is 1.5. As noise is added, the SOR threshold reduces in both cases, but it is clear that the Gaussian slice profile is less robust to the presence of noise. These results indicate that the "true" resolution

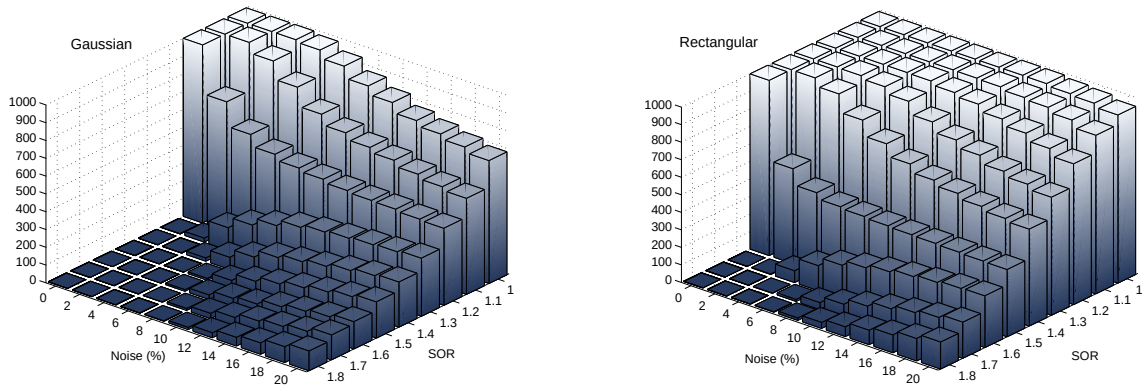

Figure 1. Histogram of successful resolutions prior to SR algorithm.

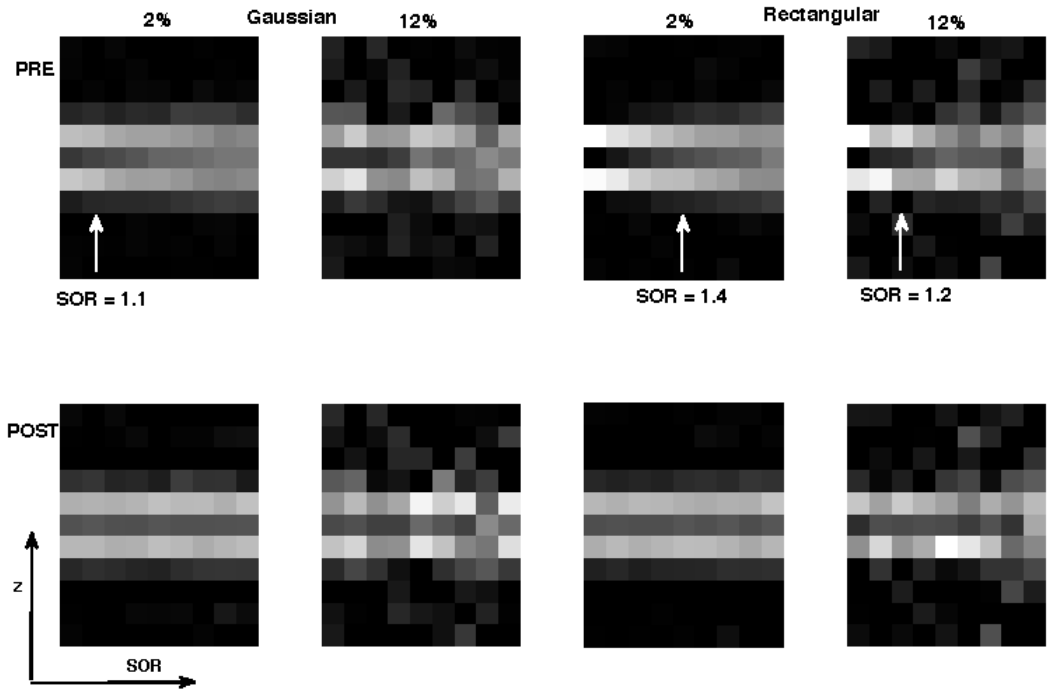

Figure 2. Examples of the simulated images (upper) and deconvolved images (lower). In all images, the vertical axis is the z-axis, the through-plane direction. The horizontal axis indicates increasing SOR. The arrows on the upper images show the maximum SOR at which the objects can be successfully resolved prior to deconvolution. (The images are zoomed to show the resolution test object only, the noise is estimated from a ROI outside of the visible area.)

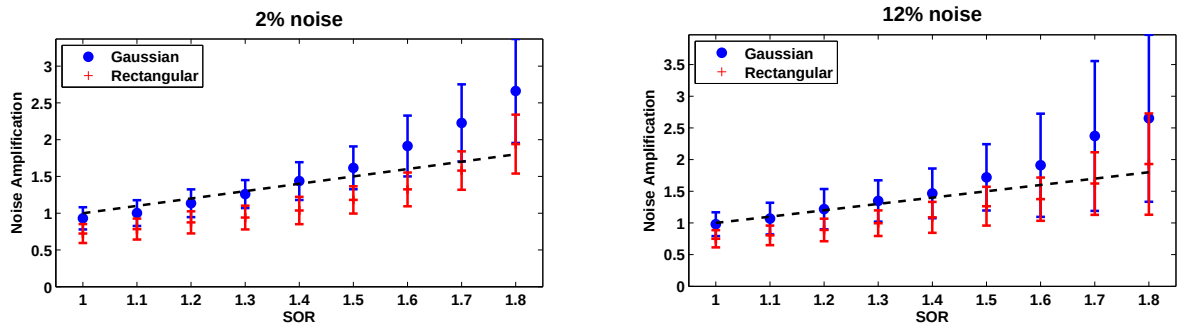

Figure 3. Graphs showing the noise amplification due to the SR algorithm, for 2% and 12% noise. The error bars indicate one standard deviation from the mean noise value. The dashed line indicates the maximum noise amplification allowed before ΔSNR falls below 1.

of the data acquired with a rectangular slice profile will often be higher than the acquired slice width suggests, even without the application of a post processing SR algorithm. Conversely, noise can mask resolution when a Gaussian slice profile is used, with the resolution test failing even for non-overlapping slices (SOR=1) at high noise levels.

SR algorithm. Selected results from the simulation and deconvolution (for noise levels 2 % and 12%) are shown in Figure 2. The upper row show the initial images prior to application of the SR algorithm, showing that the two objects are more clearly resolved, and to a higher SOR factor, when a rectangular slice profile is used. The lower row in Figure 2 show the deconvolved images from the SR algorithm. At 2% noise, deconvolution enables good visual resolution of the two objects for both the rectangular and Gaussian slice profiles. However, this resolution increase comes at the expense of noise amplification, as is illustrated by the 12% noise images in Figure 1. Figure 3 shows the noise amplification post deconvolution, for both slice profiles at different SOR factors. The Gaussian deconvolution amplifies noise more than a rectangular slice profile, and generally has a much greater variability in noise amplification, indicated by the larger error bars. This behaviour is exactly what is expected from calculating the Condition Numbers of the PSFs (data not shown). For a given SOR, the Gaussian PSF has a higher condition number compared to the rectangular PSF and hence will be more sensitive to the presence of noise, as is shown by the data in Figure 3 [10]. The dashed line in Figure 3 indicates the maximum noise amplification that can be compensated for by the increase in signal, resulting

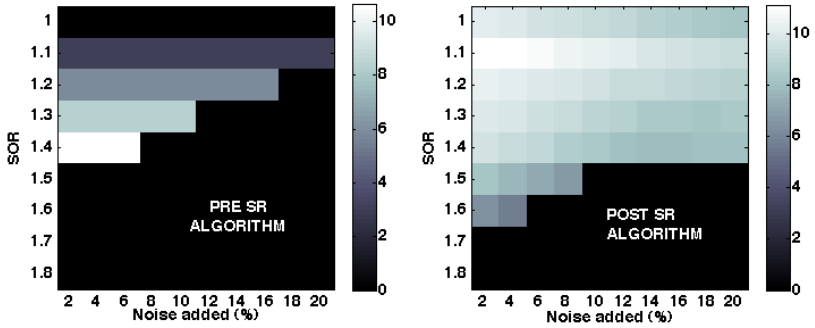

Figure 4. The mean % voxel change realised by trading increased SNR for resolution (Rectangular slice profile).

in an overall increase in SNR. For the Gaussian PSF, none of the data points and associated error bars, for all noise levels, fall below this threshold, implying that the noise amplification always exceeds the signal increase. However, for the rectangular PSF, the noise amplification is sufficiently low for a range of SOR and noise values, for example SOR ≤ 1.6 at 2% noise, and SOR ≤ 1.4 at 12% noise.

Figure 4 shows the Δx_{iso} possible from the overall ΔSNR, for the rectangular PSF, for both pre and post SR algorithm (results are thresholded to only those with 84% confidence of successful resolution). For high noise levels ($\geq 8\%$), there is a significant benefit from processing with the SR algorithm (acquiring at SOR=1.1), allowing at least a 9% reduction in voxel size for all noise levels. This improvement is greater than that expected from just the signal increase alone, suggesting that the noise has actually been reduced by the deconvolution; the result of the SR algorithm was to blur the image rather than sharpen edges. This is possible in these cases because the initial image was much sharper than required by the modulation assessment of resolution. For noise levels below this, 2-6%, roughly the same improvement can be achieved by simply acquiring at SOR=1.4 and applying no SR algorithm. This has the advantage that no deconvolution needs to be done, which requires optimisation to select the regularisation parameter. Although this simulation had an easily assessable criterion (modulation), it is harder to select and optimise such criteria with real images.

3.2 MR Test Data

Figure 5 shows an example of the results of the simulation on the MR SPGR data for SOR=1.4 and 2% noise. This figure illustrates the earlier results. Firstly, that the Gaussian slice profile results in lower resolution images than the rectangular (comparing **b** and **d**). Secondly, there is little visual difference between SOR=1 and SOR=1.4 with a rectangular slice profile (**a** and **b**), which supports the proposal that they both have resolving power equal to the matrix size. Thirdly, the noise amplification from deconvolving the Gaussian slice profile is clear to see in **e**.

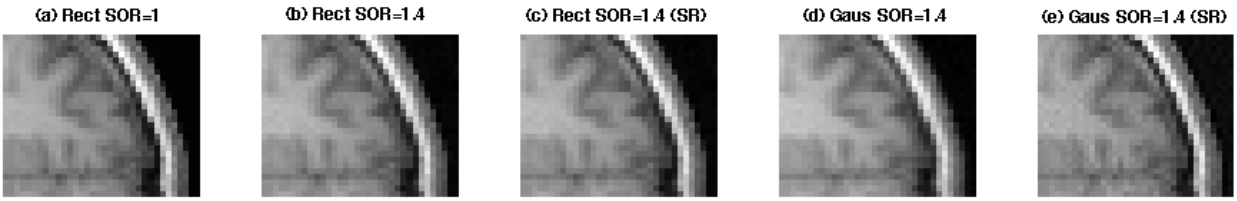

Figure 5. Examples of deconvolution on simulated MR data, with SOR=1.4 and 2% noise.

4 Conclusions

There are two main conclusions of this work. Firstly, deconvolution will always in general, come with some noise penalty. Intuitively, the deconvolution process in all SR algorithms will act similarly on noise as it does on signal, and so this is expected. The noise amplification is dependent on the slice profile size and shape, and for some situations will indeed be compensated for by the increase in signal from the SR acquisition. We have explored two different slice profiles, rectangular and Gaussian. Although rectangular profiles are the ideal slice profile, due to technological limitations this is not achievable in practice, and the real slice profile tends to be more Gaussian-like. Therefore the results from both are relevant and represent the two extremes.

We have shown that Gaussian slice profiles are not robust to the presence of noise, and the noise amplification from the deconvolution tends to outweigh the signal increase, leading to an overall reduction in SNR compared to acquiring directly at ss mm resolution. Gaussians have been shown to be one of the most ill-posed PSFs for deconvolution [10]. The situation is more hopeful for rectangular slice profiles. However, if we wish the SR algorithm to recover the exact ss mm resolution, then we have shown that there is a maximum SOR factor, which depends on the original noise in the image. At typical 2% noise, the maximum SOR is 1.6, and this reduces as noise increases. These results imply that for those studies in the literature that have quoted SOR factors above this threshold, they either carry a SNR penalty, or the SR algorithm is not truly recovering the resolution back to the desired resolution.

The second conclusion of this report relates to the initial resolution assessments, which indicate that the resolution power of a slice profile can be far higher than its slice width suggests. This has the advantage that the images need not undergo any post-processing that may amplify the noise and hence this is indeed "free" SNR. Current work attempts to use these results with real MRI data acquired with SOR>1.

This study has some practical implications. The major implication follows from the latter conclusion, and offers a free SNR boost for structural MRI (assuming that the slice profile is close to rectangular in shape). For example, an image can be said to have through-plane resolution ss mm, even when acquired at SOR=1.4 with a rectangular slice profile. This could be motivation to alter structural image acquisition parameters. Also, the dependence on slice profile shape implies that it is worth investing in improvements to the shape of the slice profile, to improve the performance of SR algorithms.

The motivation for this study was the wish to use SR to improve the resolution in SNR limited applications, such as DWI. Unfortunately, results have shown that low SNR also acts as a performance limiter for the SR algorithm. This is especially pertinent for DTI studies, as the accuracy and precision of the post-processing is very sensitive to the images SNR. For example, it has been shown that low SNR can introduce a positive bias in the Fractional Anisotropy [11]. However SR could also be used in areas where there are technical limitations to acquiring thin slices, such as gradient specifications. In this case, if the application can cope with an SNR penalty, then SR does obviously still offer potential for increasing the resolution.

Further work is intended to investigate the dependence on slice profile shape, and deconvolution routine. The simulation was carried out using a simple implementation of a SR algorithm, with deconvolution regularised by the Laplacian smoothing operator. More sophisticated SR algorithms may behave subtly differently; however, we believe the results are transferable across to other SR algorithms due to the common deconvolution that is necessary to de-blur the image. Similar results were achieved with a Wiener filter (not shown). Deconvolution algorithms used in the literature tend to be iterative based ones such as POCS and the Irani Peleg Algorithm, and it would be useful to investigate these [1].

Acknowledgements

HJ Coward acknowledges the KCL Annual Fund Studentship for funding, and Dr Bill Crum for useful discussions.

References

1. S. C. Park, M. K. Park & M. G. Kang. "Super-resolution image reconstruction: a technical overview." *Signal Processing Magazine, IEEE* **20**, pp. 21–36, 2003.
2. S. Peled & Y. Yeshurun. "Superresolution in MRI: application to human white matter fiber tract visualization by diffusion tensor imaging." *Magn Reson Med* **45**, pp. 19–35, 2001.
3. G. S. Mayer & E. R. Vrscay. "Measuring information gain for frequency-encoded super-resolution MRI." *Magnetic Resonance Imaging* **25**, pp. 1058–1069, 2006.
4. K. Scheffler. "Superresolution in MRI?" *Magn Reson Med* **48**, pp. 408, 2002.
5. H. Greenspan, G. Oz, N. Kiryati et al. "MRI inter-slice reconstruction using super-resolution." *Magnetic Resonance Imaging* **20**, pp. 437–446, 2002.
6. R. R. Peeters, P. Kornprobst, M. Niklova et al. "The use of super-resolution techniques to reduce slice thickness in functional MRI." *International Journal of Imaging Systems and Technology* **14**, pp. 131–138, 2004.
7. A. W. Dowsey, J. Keegan, M. Lerotic et al. "Motion-compensated MR valve imaging with COMB tag tracking and super-resolution enhancement." *Med Image Anal* **11**, pp. 478–91, 2007.
8. M. Bertero & P. Boccacci. *Introduction to Inverse Problems in Imaging*. Institute of Physics publishing, London, 1998.
9. R. Lerksi, J. de Wilde, D. Boyce et al. "Quality control in magnetic resonance imaging." *IPEM Report 80* 1995.
10. J. Philip. "The most ill-posed non-negative kernels in discrete deconvolution." *Inverse Problems* **3**, pp. 309–328, 1987.
11. C. Pierpaoli & P. Basser. "Toward a quantiative assessment of diffusion anisotropy." *Magn Reson Med* **36**, pp. 893–906, 1996.